

Explaining Realized Days on Market for Residential Listings in the Turkish Housing Market

Sudenas Dikici and Orhan Feyzioğlu
Department of Industrial Engineering
Galatasaray University, Istanbul, Türkiye
sudenas.dikici@ogr.gsu.edu.tr
ofeyzioglu@gsu.edu.tr

Abstract

This study examines realized Days on Market (DOM) for residential listings in Türkiye using a nationwide dataset of more than 650,000 online listings from 2024–2025. The study develops an interpretable machine-learning framework to explain observed listing duration in a heterogeneous housing market and compares a regularized linear baseline with tree-based boosting models, namely XGBoost and Histogram-based Gradient Boosting, using harmonized features related to pricing behavior, physical property characteristics, and multi-level spatial information. Model performance is evaluated using MAE, RMSE, R^2 , and relative error on held-out data. The best-performing specification, based on XGBoost, achieves held-out performance consistent with the inherent difficulty of DOM prediction ($R^2 \approx 0.39$, MAE ≈ 43 days against a mean DOM of 147 days), indicating that substantial unexplained variance remains even after accounting for pricing behavior, property attributes, and spatial heterogeneity. To move beyond aggregate model fit, the study employs gain and permutation importance, SHAP values, and partial dependence plots to interpret model behavior. The results indicate that pricing-related variables and location-based factors play a substantially greater role than most physical housing attributes in explaining realized DOM, and the divergence between gain-based and SHAP-based rankings reveals that price variables drive model fit while spatial variables dominate individual-level predictions. Overall, the findings provide a scalable and interpretable retrospective framework for diagnosing listing-side housing-market liquidity, with implications for pricing strategy, investment benchmarking, and spatial market monitoring in the Turkish housing market.

Keywords

Days on Market, Housing Market, Machine Learning, Interpretable Modeling, Pricing Strategy

1. Introduction

Housing markets are complex systems influenced by social, spatial, and behavioral factors, as well as investment and shelter needs. Property attributes, location, and listing strategy significantly affect sales outcomes together with price, supply-demand balance, and financing conditions. Prior studies have shown that listing duration and market liquidity are strongly associated with pricing behavior, seller expectations, and local market structure (Anglin et al. 2003; Haurin 1988; Glower et al. 1998; Tucker et al. 2013). These interactions are particularly visible in developing countries with uneven market liquidity, complicating decisions for buyers, sellers, investors, and policymakers.

Meanwhile, Days on Market (DOM), namely the time a property remains active on the market until the observed termination of the listing, has become a key indicator of market liquidity and efficiency. It reflects pricing and marketing effectiveness, investor capital lock-up, and overall market performance. While most studies focus on price prediction, research on listing duration remains limited, despite its importance in revealing overpricing, weak demand, or poor listing strategies.

Over the past two decades, Türkiye's housing market has changed rapidly due to urbanization, evolving credit conditions, and the growing view of housing as an investment. These shifts have altered both prices and how long properties remain on the market, creating large disparities across districts and neighborhoods. Significant differences in DOM across regions and neighborhoods highlight the roles of location, pricing, and listing tactics beyond property features alone. This study examines the determinants of housing listing duration through data-driven analysis using extensive Turkish listing data. DOM is modeled not only as a regression problem but as a means to understand market liquidity and listing behavior across various price, physical, and locational factors. The study adopts a retrospective explanatory design, using both listing-time attributes and pricing-path information observed over the life of each listing. The scope and practical value of this design are discussed in Section 3.1.

1.1 Objectives

This study investigates the factors associated with real-estate Days on Market (DOM) in Türkiye using data analytics and machine learning. In particular, the study seeks to address the following research questions:

- What are the key factors influencing DOM in the Turkish housing market?
- How do pricing behavior and physical housing characteristics affect DOM?
- How effectively can machine-learning models, particularly XGBoost, explain listing duration?
- How does hazard-based duration analysis contribute to understanding the temporal dynamics of listing duration?
- How do the results vary between nationwide and location-specific analyses?

To address these questions, the study develops interpretable machine-learning models and employs explainability techniques to examine the relative importance of the determinants of realized DOM.

2. Literature Review

This section reviews the literature on DOM and Time on Market (TOM), focusing on pricing behavior, spatial heterogeneity, machine-learning approaches, duration-based modeling, and interpretability in housing-market analysis.

In real-estate research, TOM is widely used as an indicator of market liquidity and reflects the trade-off between rapid sale and price maximization. Prior studies show that pricing behavior, seller expectations, and listing history significantly affect DOM, although empirical findings on the TOM–price relationship remain mixed across markets and periods (Lippman and McCall 1986; Anglin et al. 2003; Tucker et al. 2013).

Studies identify listing price and pricing strategy as principal determinants of DOM. Overpricing reduces buyer interest and prolongs exposure (Anglin et al. 2003), while mid-listing price adjustments both reflect seller tactics and influence sale duration (Knight 2002). Property characteristics such as size, number of rooms, age, and floor level also affect DOM, while macroeconomic conditions and local market structure contribute to variation across submarkets (Haurin 1988; Glower et al. 1998).

Early studies primarily relied on linear regression models, although such approaches often struggle to capture the nonlinear structure of housing markets (Ermolin 2016). More recent studies increasingly employ machine-learning methods for DOM analysis and real-estate decision support. For example, Castelli et al. (2020) and Cajias and Freudenreich (2024) demonstrate the usefulness of ML approaches for modeling listing duration in real-estate markets. Consequently, tree-based boosting methods have become particularly attractive for capturing nonlinear interactions and spatial heterogeneity in housing data (Chen and Guestrin 2016; Zhu et al. 2016).

Within the Turkish context, machine-learning methods have mainly been applied to housing-price prediction rather than listing duration. Existing studies focus on price forecasting and bubble detection in major metropolitan areas (Erkek et al. 2020; Ayan and Eken 2021). To the authors' knowledge, no prior study has applied machine-learning methods to explain or predict DOM in the Turkish housing market. The present study addresses this gap directly.

Housing market heterogeneity across regions significantly affects DOM. Even properties with identical features may exhibit distinct sales durations within or across urban areas (Smith, 2009). Prior studies suggest that region-specific models can better capture geographic dependencies and local market dynamics than a single unified specification (Zhu et al. 2016).

DOM can also be conceptualized as a time-to-event variable, making survival and hazard models relevant for representing listing duration as a stochastic process (Cox 1972; Lancaster 1990). Discrete-time hazard approaches are especially appropriate when duration is recorded in intervals such as days or weeks (Liu et al. 2024). In the present study, however, this perspective is used only as a conceptual reference point rather than as a fully estimated empirical benchmark.

In DOM analysis, it is important not only to assess model fit but also to understand which factors are associated with listing duration. Consequently, explainable artificial intelligence (XAI) techniques such as SHAP have gained increasing attention because they make it possible to assess how variables contribute to fitted outcomes and improve interpretability in housing-market analysis (Lundberg and Lee 2017; Molnar 2020).

Overall, the literature suggests that DOM is shaped by pricing strategy, property characteristics, market conditions, and spatial heterogeneity. However, limited research combines high-resolution location data with interpretable machine-learning analysis of DOM at a national scale in Türkiye. In this context, the present study contributes to the literature by employing a nationwide dataset, comparing global and city-specific models, and integrating multiple explainability techniques within a unified analytical framework.

3. Methodology

3.1 Scope: Retrospective Explanation vs. Listing-Time Forecasting

The modelling framework developed in this study is retrospective by design. Several features central to the analysis (notably the number of price revisions, the final observed price, and the trajectory of price adjustments) are observed over the lifetime of a listing and are therefore unavailable to a seller or platform at the moment a property is first listed. This distinction has direct implications for how the results should be interpreted and used.

The practical value of a retrospective framework lies not in real-time listing-time prediction but in three analytically distinct uses:

- **market diagnostics:** identifying which pricing and locational patterns are systematically associated with prolonged DOM at an aggregate level, informing platform analytics and market monitoring.
- **investment and policy benchmarking:** allowing investors, developers, and public authorities to assess historical liquidity conditions across submarkets and price segments.
- **indirect seller guidance:** the finding that initial listing price and pricing revision frequency are the strongest retrospective correlates of DOM provides indirect evidence that accurate initial pricing reduces expected duration more reliably than subsequent price adjustments, even though this inference cannot be read as strictly causal from the present design.

A strictly listing-time model, one using only features observable at the moment of listing, would be more directly actionable for real-time decision support but is necessarily less informative about DOM dynamics, because it excludes the pricing evolution that turns out to be the dominant correlate of listing duration. Both perspectives are methodologically legitimate; the present study occupies the retrospective pole and should be read accordingly. An ex ante extension of this framework, incorporating only listing-time features, is identified as a priority direction for future work in Section 6.

3.2 Modelling Approach

This study employs an integrated methodological framework to analyze realized DOM in the Turkish housing market and to identify the factors most strongly associated with listing duration. Machine-learning models are used to capture nonlinear relationships and interactions, while interpretability tools are employed to examine the relative importance of the key covariates.

The empirical framework combines listing-time attributes with observed pricing-history information, including initial price, final observed price, and the number of price revisions during the listing period. Physical housing features include area, number of rooms, building age, floor level, and total number of floors, while location variables capture spatial variation at provincial, district, and neighborhood levels.

Because DOM involves nonlinear relationships and substantial spatial heterogeneity, the study compares multiple modeling families within a unified framework. Ridge Regression is used as a regularized linear baseline, whereas the

primary analytical focus lies on XGBoost and Histogram-based Gradient Boosting (HistGB), both of which are well suited to capturing complex interactions in tabular housing data. Hyperparameters were tuned using K-fold cross-validation on the training set, and configurations were selected according to validation RMSE and overfitting control.

Model performance is evaluated using MAE, RMSE, and R^2 metrics together with residual analysis in order to assess predictive accuracy and model stability across the DOM distribution. To examine spatial heterogeneity, both a nationwide global model and city-specific submodels are estimated for major metropolitan areas.

In addition to the regression-based analysis, the study acknowledges a discrete-time hazard perspective that conceptualizes DOM as a stochastic duration process. In the present work, however, this perspective is included only to contextualize the temporal interpretation of DOM and to motivate future methodological extensions rather than as a fully estimated empirical benchmark.

4. Data, Experimental Design, and Evaluation Framework

The dataset employed in this research comprises extensive residential listing data sourced from online real estate platforms throughout Türkiye. It includes details on prices, property characteristics, and locations that could impact listing duration or DOM. Aligning with the study's primary objective, the data collection and preprocessing aimed to enhance model accuracy while establishing a data framework for analyzing how listing duration relates to economic and spatial factors.

4.1 Data Source and Scope

The dataset utilized in this study was obtained from a major online real-estate platform operating in Türkiye, with each record representing a single residential property listing. Covering listings from 2024 to 2025, the dataset includes information such as listing dates, observed listing end dates, price revisions during the listing period, physical property characteristics, and multi-level location attributes (province, district, and neighborhood). Data preparation and preprocessing procedures were conducted to construct a structured analytical dataset suitable for large-scale empirical modeling. No personally identifiable user information was included in the analysis.

Accordingly, DOM is interpreted throughout the paper as listing duration up to the observed termination of the listing, rather than as confirmed time-to-sale only. A limitation of this definition is that listing termination does not necessarily imply a confirmed transaction. Some listings may have been withdrawn, expired, temporarily suspended, or otherwise removed from the platform without resulting in an actual sale. Consequently, the observed DOM variable should be interpreted as realized listing duration rather than verified time-to-sale. Despite this limitation, listing termination remains a meaningful proxy for market liquidity and seller-side listing outcomes in large-scale online housing-market analyses.

Its extensive geographic coverage and large number of observations facilitate analysis of the Turkish housing market's heterogeneity. The main variable, DOM, was calculated as the total number of active days observed for each listing in the retained sample. Given its right-skewed distribution, a logarithmic transformation was applied to stabilize variance and minimize outlier effects, and models were trained on $\log_{1p}(\text{DOM})$. Despite this transformation, we will show that prediction errors remain larger for very long marketing durations, reflecting residual skewness and heterogeneity in the tail of the DOM distribution. During preprocessing, missing or inconsistent data were checked, outliers were constrained using domain-based thresholds, and a clean, ready-to-model dataset was prepared. As shown in Table 1, some variables are not available for all listings, which leads to variable-specific differences in descriptive sample counts.

Table 1. Descriptive statistics of the cleaned housing dataset

Variable	Column	N	Mean	Median	Std	Min	Max
Total observations	–	658,416					
Number of provinces	Province	81					
Number of districts	District	747					
Number of neighborhoods	Neighborhood	7,303					
Days on Market	DOM	658,416	150	109	92	1	747

Initial price	InitialPrice	658,416	4,672,725	3,100,000	5,168,556	50,000	200,000,000
Final price	FinalPrice	658,416	4,701,283	3,100,000	5,161,346	50,000	100,000,000
Price revisions	PriceRevisions	658,416	2	1	1	1	54
Net area (m ²)	NetArea	658,416	119	110	54	10	500
Gross area (m ²)	GrossArea	658,416	139	125	70	11	800
Building age	BuildingAge	658,416	11	6	11	0	80
Floor level	FloorLevel	418,424	3	3	3	1	21
Total floors	TotalFloors	658,416	5	4	4	1	60
Number of rooms	NumberofRooms	658,416	3	3	1	1	10
Price per m ² (TL/m ²)	PricePerM2	658,416	39,210	29,077	30,537	200	200,000

4.2 Variables and Feature Groups

The variables in the dataset are categorized into three main groups to maintain analytical clarity:

- **Price Dynamics:** Initial price, final observed price, and the number of price adjustments recorded during the listing period. In the present study, these variables are interpreted as ex post indicators of pricing behavior over the life of the listing rather than as information available at the initial listing moment. Note that PriceRevisions refers to the number of observed price records over the life of the listing, so the initial posted price is counted as the first revision.
- **Physical Housing Features:** Variables include net and gross square meters, number of rooms, building age, floor level, total floors, heating type, housing type, and whether the property is part of a housing complex.
- **Location Information:** Detailed categorical location variables are available at the province, district, and neighborhood levels. This structure is maintained in order to incorporate spatial heterogeneity into the explanatory framework.

Numerical variables were scaled prior to modeling, while categorical variables were converted using one-hot encoding. This approach resulted in a unified feature space suitable for both linear and tree-based models. We deliberately focused on primitive listing-level features and did not include more elaborate engineered variables such as local average DOM or relative pricing gaps. These are discussed as potential extensions in Section 6.

4.3 Data Cleaning and Preprocessing Process

Observations with logical inconsistencies were excluded, while missingness was handled during preprocessing in line with the modeling framework. Outliers in DOM and price variables were constrained using domain-based thresholds in order to reduce the influence of extreme observations on model training. Because the target variable exhibited a right-skewed distribution, a log1p transformation was applied prior to modeling. The resulting dataset is suitable for both linear and nonlinear machine-learning models.

4.4 Training–Test Split and Modeling Framework

The dataset was divided into training, validation, and test subsets. Hyperparameter tuning was conducted using K-fold cross-validation within the training portion, and final model performance was reported on the held-out test set. Table 3 therefore reports both validation-stage and final test-set performance. We also verify that DOM distributions in the training and test sets are aligned, ensuring that performance metrics are not driven by distributional shifts.

In the study, a log1p transformation was used to improve model stability and reduce the influence of extreme values, as the target variable DOM exhibited a right-skewed distribution. All models were trained on this transformed target. XGBoost serves as the primary nonlinear explanatory model because of its ability to capture complex DOM structures and intricate variable interactions. Ridge regression serves as a linear baseline for comparison, while the Histogram-based Gradient Boosting (HistGB) model is incorporated as an alternative tree-based method for evaluating robustness.

The analysis used 16 raw variables; after one-hot encoding of categorical variables, the effective feature space expanded substantially. Model hyperparameters were adjusted within the recommended ranges reported in the

literature. The final model configurations are presented in Table 2. For the HistGB model, the 'Trees' parameter in the table corresponds to the max_iter setting in scikit-learn library.

Table 2. Models used and basic hyperparameter settings

Model	Seed	Features	Dataset	Trees	Depth	Learning Rate	Subsample	Colsample	Alpha
Ridge	42	16	Clean dataset	–	–	–	–	–	1
XGBoost	42	16	Clean dataset	2,200	8	0.03	0.80	0.80	–
HistGB	42	16	Clean dataset	600	–	0.05	–	–	–

4.5 Performance Criteria

Model performance was assessed using standard regression metrics that capture both error magnitude and explanatory fit. The primary metrics are Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), coefficient of determination (R^2), and MAE relative to mean DOM. These metrics permit comparison of the models in terms of held-out fit and practical interpretability. Notably, the MAE/Mean DOM ratio provides an operational indication of the magnitude of model error relative to the average observed listing duration. Table 3 presents the results for the validation and test datasets.

Table 3. Performance of the models on validation and test datasets

Model	Set	MAE	RMSE	MSE	R^2	MAPE	MAE /Mean DOM	DOM		
								Mean	Median	Std. Dev.
Ridge (log1p target)	VALID	51.26	164.40	27027.88	-2.259	29.63	0.34	146.96	105	91.07
Ridge (log1p target)	TEST	51.62	151.21	22865.79	-1.708	29.81	0.35	147.43	105	91.89
XGBoost (log1p target)	VALID	43.35	71.39	5096.79	0.386	25.71	0.29	146.96	105	91.07
XGBoost (log1p target)	TEST	43.41	71.51	5114.15	0.394	25.63	0.29	147.43	105	91.89
HistGB (log1p target)	VALID	43.89	72.25	5220.46	0.371	25.95	0.29	146.96	105	91.07
HistGB (log1p target)	TEST	44.02	72.51	5258.85	0.377	25.88	0.29	147.43	105	91.89

The results in Table 3 show that the XGBoost model achieves the lowest error metrics on the test set, with an MAE of 43.41 days and an RMSE of 71.51 days. This corresponds to roughly a 29.4% relative error based on the average observed listing duration. While the HistGB model shows comparable performance with an MAE of 44.02 days, XGBoost performs slightly better in both MAE and R^2 .

The Ridge regression model, constrained by its linear structure, is less able to capture the nonlinear relationships associated with DOM and therefore yields substantially weaker fit than the tree-based models. Its negative R^2 values indicate that the linear specification performs worse than a constant-mean benchmark. The close alignment between validation and test results for XGBoost and HistGB suggests that overfitting is limited. Based on these observations, XGBoost was selected as the primary model for the subsequent explanatory and interpretive analyses.

4.6 Regional and Spatial Heterogeneity

It is well-known that the Turkish housing market exhibits notable spatial heterogeneity. To evaluate how this heterogeneity affects model performance, city-specific XGBoost models were created alongside the global model. Specifically, separate models were trained for Ankara, Istanbul, and İzmir, and their performances were compared to that of the global model. Table 4 provides a side-by-side comparison of the global and city-based models' performance on the test dataset, using R^2 and MAE metrics.

Table 4. Comparison of global and city-specific model performance

City	R ² _Global	R ² _City	ΔR^2 (City-Global)	MAE_Global	MAE_City	ΔMAE (City-Global)	N_test
Ankara	0.394	0.404	+0.009	43.42	32.08	-11.33	25,004
İstanbul	0.394	0.371	-0.023	43.42	43.50	+0.08	21,851
İzmir	0.394	0.325	-0.069	43.42	54.97	+11.55	7,310

The results reveal that spatial heterogeneity affects model performance differently across submarkets. The Ankara-based model exhibits a modest improvement over the global model ($\Delta R^2 \approx 0.009$ and $\Delta MAE \approx -11$ days), indicating some gains from local calibration in a relatively homogeneous submarket. By contrast, the city-specific models for Istanbul and İzmir perform similarly to or worse than the global specification. This suggests that, in these cases, greater intra-city heterogeneity and more complex submarket structure may outweigh the benefits of separate calibration. In Istanbul, the R² value declines slightly while MAE remains nearly unchanged. In İzmir, the deterioration is more pronounced, with both lower R² and higher MAE. Overall, these findings indicate that spatial segmentation does not uniformly improve held-out model performance.

4.7 Outputs Used for Model Interpretability

Model interpretability was evaluated using gain-based feature importance, permutation importance, SHAP values, and partial dependence plots (PDPs). Gain and permutation importance identify variables that contribute most strongly to predictive fit and held-out performance, whereas SHAP values provide observation-level explanations of fitted outcomes. PDP analyses were additionally used to examine the marginal relationships between selected numerical variables and log(DOM). The resulting interpretability findings are presented in Sections 5.3 and 5.4.

5. Results and Discussion

5.1 Numerical Results

This section reports and interprets the numerical results of the developed models. Table 3 shows that the XGBoost model achieves the strongest held-out performance among the evaluated specifications, with a test-set MAE of 43.41 days and an R² of 0.394. The HistGB model produces comparable results, whereas Ridge regression performs substantially worse because its linear structure cannot adequately capture the nonlinear relationships associated with DOM. Validation and test results are closely aligned for the tree-based models, indicating limited overfitting.

Although the XGBoost model provides the strongest performance among the evaluated specifications, the resulting R² value of approximately 0.39 indicates that a substantial proportion of variation in realized DOM remains unexplained. This level of explanatory power is nevertheless comparable to previous DOM studies and reflects the inherently complex and partially unobservable nature of housing markets. Factors such as seller motivation, negotiation processes, buyer preferences, financing conditions, and local market dynamics are difficult to observe directly and therefore cannot be fully captured within the current framework. Consequently, the primary contribution of the proposed approach lies not in maximizing predictive accuracy, but in providing an interpretable framework for understanding the key determinants of residential listing duration.

To examine spatial heterogeneity, the global specification was compared with city-specific XGBoost models for Ankara, Istanbul, and İzmir (Table 4). The Ankara model performs slightly better than the global specification, whereas the Istanbul and İzmir models perform similarly or worse. These findings suggest that greater intra-city heterogeneity may limit the benefits of separate local calibration. Table 5 reports the districts with the shortest observed listing durations in each city.

Table 5. Districts with the shortest listing durations

City	District	N	DOM Mean	DOM Median	%DOM ≤ 90	%DOM ≤ 180	Rank
Ankara	Sincan	20,736	116.13	93	41.31	91.38	1
Ankara	Keçiören	32,527	118.56	96	36.30	90.68	2
Ankara	Pursaklar	12,034	120.46	96	40.16	88.37	3

Ankara	Etimesgut	17,950	129.15	101	31.67	85.06	4
Ankara	Çankaya	38,226	138.11	102	31.34	81.33	5
İstanbul	Esenler	2,256	125.43	90	53.50	85.77	1
İstanbul	Bağcılar	2,886	132.38	90	54.75	82.81	2
İstanbul	Bayrampaşa	1,141	139.26	90	54.34	79.75	3
İstanbul	Beşiktaş	5,396	155.79	90	50.09	74.57	4
İstanbul	Esenyurt	10,514	126.61	91	48.69	86.32	5
İzmir	Torbalı	1,873	144.03	102	39.24	79.12	1
İzmir	Narlıdere	454	156.60	105	39.21	73.35	2
İzmir	Menemen	3,039	144.12	106	39.16	77.72	3
İzmir	Buca	4,229	146.61	106	36.27	77.11	4
İzmir	Bayraklı	1,577	153.23	109	38.81	73.68	5

The shortest observed listing durations in Ankara are found in Sincan, Keçiören, and Pursaklar. In Istanbul, Esenler, Bağcılar, and Bayrampaşa exhibit median DOM values of approximately 90 days, indicating relatively high listing-side liquidity. In İzmir, Torbalı and Menemen display the shortest listing durations within the city. Overall, the findings indicate that residential listing duration is strongly associated with local submarket conditions in addition to pricing behavior and property characteristics.

5.2 Hazard-Based Duration Analysis

To complement the regression-based framework, a discrete-time hazard model was implemented to examine the temporal dynamics of listing closure. Unlike regression models, which estimate a single expected DOM value, the hazard framework focuses on the conditional probability that a listing closes during a particular interval.

The estimated baseline hazard pattern indicates that closure probabilities vary across time, suggesting that housing market activity follows a dynamic rather than static process. The probability of listing closure changes throughout the marketing period, implying that DOM should be interpreted as a time-dependent phenomenon rather than a fixed outcome. These findings support the view that duration analysis provides complementary insights beyond standard machine-learning predictions.

Overall, the hazard-based framework and the machine-learning models answer different but related questions. While XGBoost focuses on explaining and predicting the expected duration of listings, the hazard approach captures how the probability of listing closure evolves over time. Therefore, both perspectives contribute to a more comprehensive understanding of housing market liquidity.

5.3 Graphical Findings

Figure 1 presents residual diagnostics for the XGBoost model on the test set. The left panel depicts residuals as a function of the observed DOM values, where residuals are defined as True – Fitted DOM. At low and moderate DOM levels, the residuals remain relatively concentrated around zero, indicating stable fitted performance in the central part of the distribution. However, prediction errors become larger as observed DOM increases, suggesting that the model has greater difficulty representing listings with unusually long market durations. Overall, the figure indicates that the model performs reasonably well for typical listings but is less accurate for extreme high-DOM observations.

The right panel summarizes the distribution of residuals for the test set. The residuals are centered close to zero, indicating limited overall prediction bias. At the same time, the distribution exhibits a left-skewed tail, suggesting that the model tends to underestimate a subset of listings with relatively long DOM values (Figure 2).




Figure 1. Residual diagnostics for the XGBoost model: residuals versus true DOM (left panel) and distribution of residuals (right panel)




Figure 2. Baseline hazard curve of housing listings

Figure 3 summarizes feature importance using both gain-based and permutation importance measures. The left panel reports gain-based importance, ranking variables according to their contribution to reducing model error, whereas the right panel presents permutation importance based on the increase in MAE when individual features are randomly perturbed. Both approaches consistently identify final observed price, the number of price changes, and the initial listing price as the dominant predictors of DOM. Location-related variables, particularly district and city, exhibit a moderate level of importance, suggesting that spatial context and submarket conditions also meaningfully shape listing duration. By contrast, physical housing characteristics such as gross and net area, building age, floor level, and total number of floors contribute comparatively less to explanatory fit. Taken together, these results indicate that pricing behavior and location are more strongly associated with listing duration than detailed structural attributes.




Figure 3. Feature importance measures for the XGBoost model: gain-based importance (left panel) and permutation importance based on MAE increase (right panel)

Figure 4 presents partial dependence plots for selected features, showing the marginal effects of final price, initial price, and net area on $\log(\text{DOM})$. The initial price displays a clear positive association with $\log(\text{DOM})$: as the starting price rises, the expected marketing period lengthens, indicating that setting an excessively high initial price tends to prolong listing duration. The effect of the final price is comparatively modest and less systematic, suggesting that, conditional on other factors, closing prices play a smaller role in determining time on the market than the initial listing price. By contrast, net area exhibits a slight negative relationship with $\log(\text{DOM})$, implying that larger dwellings have shorter observed DOM on average. Taken together, these patterns point to nonlinear relationships between key covariates and DOM and underscore that the initial price is an important explanatory variable associated with the duration of the listing process.




Figure 4. Partial dependence plots for selected features on $\log(\text{DOM})$ (XGBoost)

Figure 5 reports the global feature importance derived from SHAP values for the XGBoost model. Spatial variables, particularly district, neighborhood, and city, emerge as the strongest contributors to individual-level DOM predictions. By contrast, price-related and traditional physical housing variables contribute relatively less once detailed location information is incorporated into the model.




Figure 5. Global SHAP feature importance (XGBoost, grouped top 20)

5.4 Reconciling Gain-Based and SHAP-Based Importance Rankings

The gain-based and permutation-based importance analyses (Figure 3) consistently rank price-related variables (final observed price, number of price revisions, and initial listing price) as the dominant predictors of DOM. The SHAP analysis (Figure 5), however, ranks spatial variables (district, neighborhood, and city) as the most influential contributors on an absolute scale, with price variables occupying only a moderate position. These two sets of results describe the same model but answer different interpretive questions.

The divergence reflects a well-established distinction between what each importance framework measures. Gain-based importance captures the total reduction in prediction error achieved by each variable across model splits, while permutation importance evaluates how much model performance deteriorates when a variable is randomly disrupted.

Because price variables are continuous and numerically granular, they are highly efficient split candidates and therefore generate large cumulative gain. SHAP importance, however, measures the average magnitude of each variable's contribution to individual predictions across the test set. Spatial variables may therefore appear more influential in SHAP-based rankings because they capture broader submarket-level shifts in predicted DOM.

Taken together, the two perspectives are complementary rather than contradictory. The model learns much of its predictive structure from price variation, but at the individual listing level, spatial embedding also produces substantial shifts in expected duration. A complete understanding of DOM in the Turkish housing market therefore requires both lenses: pricing behavior explains how liquidity varies within a submarket, while location helps explain why liquidity differs across submarkets.

5.5 Limitations and Future Refinement

The present results indicate that the proposed framework captures the overall structure of the data and provides useful explanatory insights into residential listing duration. At the same time, several limitations should be considered when interpreting the findings. In particular, prediction errors become larger for listings with unusually long DOM values, suggesting that the upper tail of the distribution remains more difficult to represent within the current specification. The retrospective design of the framework also implies that the practical value of the analysis lies primarily in market diagnostics and investment benchmarking rather than in real-time listing-time forecasting.

The interpretability analyses consistently show that pricing-related variables play a central role in explaining model behavior, while the strong contribution of spatial variables highlights the importance of submarket heterogeneity in the Turkish housing market. Although the present framework captures this heterogeneity through detailed location variables and city-level comparisons, future work may explore richer spatial representations, alternative feature constructions, and ex ante modelling approaches using only listing-time information.

6. Conclusion

This study investigates the determinants of DOM for residential listings in the Turkish housing market using interpretable machine-learning methods and large-scale listing data. Across the evaluated model families, XGBoost delivers the strongest held-out performance, with an average error of approximately 29% relative to mean observed DOM. The results indicate that pricing behavior and spatial heterogeneity are the primary factors associated with realized listing duration.

The interpretability analyses show that both pricing dynamics and spatial variables play major roles in explaining DOM. Gain-based and permutation-based measures emphasize the importance of pricing behavior, whereas SHAP-based explanations highlight the strong contribution of spatial context to individual-level predictions. The findings also reveal substantial spatial heterogeneity across Turkish housing submarkets, although city-specific models do not consistently outperform the global specification.

The study contributes to the literature by providing one of the few large-scale empirical analyses of realized DOM in the Turkish housing market and by integrating multiple explainability techniques within a unified analytical framework. Beyond aggregate model performance, the analysis demonstrates how pricing behavior, location, and nonlinear interactions are associated with observed listing duration in a spatially diverse housing market.

The primary objective of this study is not to maximize predictive accuracy, but to provide an interpretable analytical framework for understanding the determinants of residential listing duration. Accordingly, the emphasis is placed on explainability and economic interpretation rather than solely on incremental improvements in predictive performance. Explainable machine-learning techniques make it possible to identify the relative importance of pricing behavior, physical housing characteristics, and spatial factors, thereby providing insights that are potentially more valuable for market analysis and decision-making than marginal gains in forecasting accuracy.

From a practical perspective, the findings may assist real-estate professionals, developers, investors, and online listing platforms in understanding the factors associated with prolonged marketing periods. The results suggest that appropriate initial pricing is more effective than repeated price revisions in improving market liquidity. In addition, the strong influence of spatial variables highlights the importance of location-specific market conditions, which may help market participants benchmark liquidity across submarkets. Online platforms may also employ similar analytical

frameworks to identify inactive listings and support pricing recommendations, thereby improving market transparency and decision-making.

At the same time, the retrospective nature of the framework should be considered when interpreting the findings. Because several explanatory variables, particularly pricing revisions and final observed prices, are only observable after the listing process unfolds, the present framework is more appropriate for retrospective market diagnostics and liquidity analysis than for real-time forecasting at the moment of listing. Accordingly, the findings should be interpreted primarily as explanatory insights into realized market behavior rather than as purely ex ante predictive tools.

Overall, the study demonstrates that interpretable machine-learning frameworks can improve understanding of realized residential listing duration in a heterogeneous housing market. Future research may extend the framework through richer feature engineering and fully estimated duration-based models, including discrete-time hazard specifications.

References

- Alves, C. J., *Airbnb price prediction in Lisbon and Porto: A machine learning approach with multi-feature integration*, Master's Thesis, Universidade NOVA de Lisboa, Portugal, 2024.
- Anglin, P. M., Rutherford, R. and Springer, T. M., The trade-off between the selling price of residential properties and time-on-the-market: The impact of price setting, *The Journal of Real Estate Finance and Economics*, vol. 26, no. 1, pp. 95–111, 2003.
- An, Z., Cheng, P., Lin, Z. and Liu, Y., How do market conditions impact price–time-on-market relationship? Evidence from real estate owned (REO) sales, *Journal of Housing Economics*, vol. 22, no. 3, pp. 250–263, 2013.
- Ayan, E. and Eken, S., Detection of price bubbles in Istanbul housing market using LSTM autoencoders: a district-based approach, *Soft Computing*, vol. 25, pp. 7895–7906, 2021.
- Cajias, M. and Freudenreich A., What are tenants demanding the most? A machine learning approach for the prediction of time on market, *Journal of Property Investment & Finance*, vol. 42, no. 2, pp. 151–165, 2024.
- Chen, T. and Guestrin, C., XGBoost: A scalable tree boosting system, *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 785–794, San Francisco, USA, 2016.
- Cox, D. R., Regression models and life-tables, *Journal of the Royal Statistical Society: Series B*, vol. 34, no. 2, pp. 187–220, 1972.
- Erkek, M., Çayırılı, K. and Hepşen, A., Predicting house prices in Turkey by using machine learning algorithms, *Journal of Statistical and Econometric Methods*, vol. 9, no. 4, pp. 29–42, 2020.
- Ermolin, S. V., *Predicting days-on-market for residential real estate sales*, CS229 Course Project Report, Stanford University, USA, 2016.
- Glomer, M., Haurin, D. and Hendershott, P., Selling time and selling price: the influence of seller motivation, *Real Estate Economics*, vol. 26, no. 4, 719–740, 1998.
- Haurin, D. R., The duration of marketing time of residential housing, *Real Estate Economics*, vol. 16, no. 4, pp. 396–410, 1988.
- Knight, J. R., Listing price, time on market, and ultimate selling price: Causes and effects of listing price changes, *Real Estate Economics*, vol. 30, no. 2, pp. 213–237, 2002.
- Lancaster, T., *The Econometric Analysis of Transition Data*, Cambridge University Press, Cambridge, UK, 1990.
- Lazear, E. P., Retail pricing and clearance sales, *American Economic Review*, vol. 76, no. 1, pp. 14–32, 1986.
- Lippman, S. A. and McCall, J. J., An operational measure of liquidity, *American Economic Review*, vol. 76, no. 1, pp. 43–55, 1986.
- Liu, Y., Shakib, S., Miller, E. J. and Habib, K. N., Modeling home property listings' time-on-market duration and listing outcome using copula-based competing risk method, *Journal of Transport and Land Use*, vol. 17, no. 1, pp. 579–601, 2024.
- Lundberg, S.M. and Lee, S.-I., A unified approach to interpreting model predictions, *Advances in Neural Information Processing Systems*, vol. 30, pp. 4765–4774, 2017.
- Molnar, C., *Interpretable machine learning: A guide for making black box models explainable*, Independently Published, 2020. Available at: <https://christophm.github.io/interpretable-ml-book/>
- Trindade Neves, F., Aparicio, M. and de Castro Neto, M., The impacts of open data and explainable AI on real estate price predictions in smart cities, *Applied Sciences*, vol. 14, no. 5, 2209, 2024.
- Smith, B.C., Spatial Heterogeneity in Listing Duration: The Influence of Relative Location to Marketability, *Journal of Housing Research*, vol. 18, no. 2, 151–171, 2009.

- Sirmans, G. S., Macpherson, D. A. and Zietz, E. N., The composition of hedonic pricing models, *Journal of Real Estate Literature*, vol. 13, no. 1, pp. 3–43, 2005.
- Suresh, K., Severn, C. and Ghosh, D., Survival prediction models: An introduction to discrete-time methods and assessment of predictive performance, *BMC Medical Research Methodology*, vol. 22, no. 1, 207, 2022.
- Tucker, C. E., Zhang, J. and Zhu, T., Days on market and home sales, *RAND Journal of Economics*, vol. 44, no. 2, pp. 337–360, 2013.
- Castelli, M., Dobрева, M., Henriques, R. and Vanneschi, L., Predicting days on market to optimize real estate sales strategy, *Complexity*, no. 1, 4603190, 2020.
- Zhu, H., Xiong, H., Tang, F., Liu, Q., Ge, Y., Chen, E. and Fu, Y., Days on market: Measuring liquidity in real estate markets, *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 393–402, San Francisco, USA, 2016.

Biographies

Orhan Feyzioğlu is a Professor of Industrial Engineering at Galatasaray University, Istanbul. He received his B.Sc. in Industrial Engineering from Istanbul Technical University and his M.Sc. and Ph.D. in Industrial Engineering from Boğaziçi University. His research interests include optimization theory and methods, supply chain and logistics management, decision theory, fuzzy sets and systems, nonlinear programming, management information systems, and quantitative methods. He has also held major academic leadership roles at Galatasaray University, including Head of Smart Systems Engineering, Director of the Technology Transfer Center, and University Executive Board Member.

Sudenas Dikici is currently pursuing a master's degree in Industrial Engineering at Galatasaray University. Her research interests include data analytics, machine learning, housing market analysis, and sustainable transportation systems. In her work, she particularly focuses on the sales duration in the real estate market (Days on Market), examining regression, classification, and survival analysis methods in a comparative way. She also conducts data-driven analyses on price dynamics, location-based features, and market liquidity. In addition to her academic studies, she has hands-on experience in process improvement and data analytics.